# Generating Light Curves from Forced PSF-fit Photometry on PTFIDE Difference-images

**F. Masci, S. Papadogiannakis, R. Laher, L. Yan, Y. Cao, & B. Cenko**

**8/11/2015, vsn 1.5**

## Outline

## 1 Introduction

This document describes the fundamentals of generating a light curve (in calibrated magnitude units versus time) from the raw output tables produced by IPAC's differential forced-photometry tool. Caveats, warnings, and suggestions for optimizing and vetting light curve measurements (for maximal S/N) are also given. This document does not tell you how to submit a forced photometry request.

The intent of forced photometry is twofold:

(i) Obtain publication-quality light curves with properly vetted uncertainties and upper limits on non-detections;

(ii) Enable deeper detection or place tighter constraints on non-detections by optimally combining noisy measurements, i.e., below the single-exposure limit.

The applications include both variables of any type and transients. For suspect periodic variables, one can use a prior period (or trial a range) to create phase-folded averaged time-series using all the forced-photometry measurements. This yields high S/N phase-folded light curves, enabling discovery at fainter regimes. Therefore, to sum up, forced photometry provides "icing on the cake". The flavor of the icing is up to you!

In contrast, the Catalog Photometry from the production pipeline represents “unforced” photometry, i.e., depends on an extraction threshold being satisfied to assist with initial discovery and quick-look photometry to guide further follow-up. This photometry is also referred to as “DC” or absolute photometry. Forced photometry performed on difference-images is sometimes referred to as “AC”, differential, or relative photometry. This can be converted to DC photometry (for example, for variables) using some estimate of the time-averaged flux-level. See Section 9 for details.

Before delving into the products, or even before submitting a forced photometry request, we encourage you familiarize yourself with Sections 2 – 10 below. Also, be prepared to write software to analyze the products and generate light curves. Every submitted request will be different, and the parameters you choose for the suggested corrections and thresholds will need visual guidance. Some of the details will not be new to you. It is presented here as a guide. Again, you’re free to decorate and slice your cake however you like.

## 2 Before you submit your request

Here are some advisories to avoid potential distress for both you and the poor machine(s) that will be processing your request. Some of the details are expanded further below.

(a) Ensure your targets (candidates) of interest have been vetted by other means, e.g., check if ancillary follow-up observations are available. Spectroscopy would be ideal, but independent confirmation is recommended. To get to the point, make sure it has a high likelihood of being real. Goes without saying, right? The forced-photometry service is not a vetting service.

(b) Ensure your supplied R.A., Dec. position(s) are correct and reflect what would be measured directly off IPAC’s astrometrically-calibrated science images in the archive. **Note:** all our astrometry pertains to Equinox J2000.

(c) Ensure you request a sufficient number of “historical” measurements prior to the epoch(s) defining your transient event or any characteristic behavior/pattern sought for. We suggest ≥ 30 epochs prior to your desired time-span of interest. This requirement is not necessary for continuous (a)periodic variables. For details, see Section 2. Therefore, it is assumed you understand your target(s) well enough to judiciously select the time-span(s).

Stay tuned for more advisories as we learn them – primarily from you.

# 3 Outputs: plot of flux [DN] versus time [MJD]: baseline correction

We recommend you always use the PSF-fitted differential fluxes and corresponding uncertainties in DN when generating an initial light curve. These are the “flux” and “sigflux” columns in the forced-photometry output tables (Appendix I). Fluxes and uncertainties from forced *fixed-aperture* photometry are also provided (“fluxap” and “sigfluxap”). These are purely ancillary and their potential use is discussed in Section 5. The PSF-fitted fluxes are generally more accurate (in terms of S/N).

It is recommended you examine a plot of the flux [DN] vs time [MJD] measurements to determine if there is any residual offset in the historical baseline. The definition of “historical” is completely up to you. For example, if the reference image used to generate the difference images (from which forced photometry is derived) was contaminated by the transient flux that’s sought, your historical baseline will be < 0. If the reference image was affected by some other systematic in the generation or calibration process, the baseline could be > 0. **Note:** requesting a special time range for the reference-image for difference-image creation to support archival forced photometry is superfluous. You will always need to correct for a possible non-zero baseline *for transients* to obtain the best photometric accuracy. This *is not necessary for variables* that may fluctuate about a constant long-term flux-level. The reason is that even if the reference image were not a true (unbiased) time-average of all the input epochs (at random light curve phases), any residual baseline is a DC level you may want to retain when converting from AC to DC photometry (see Section 9).

The baseline correction needs to use a sufficient number of historical measurements. We suggest ≥ 30 epochs. Therefore, it is imperative that you have some prior knowledge of the behavior of your transient and go back far enough to include a sufficient historical set of epochs before submitting your request. Figure 1 shows an example of a transient and the flux (DN) vs time plot used to derive the baseline correction. Raw data is shown on the left and the corrected measurements (after subtracting 35 DN) are on the right. Note that the time-range used to compute this correction is entirely up to you. This could be driven by your science application. For example, Figure 1 shows some early activity starting at $\mathrm{MJD} - \mathrm{MJD}_0 \sim -45$ days (relative to the peak epoch). Therefore, to avoid any bias from this activity, you might want to compute the baseline correction using only measurements at ≤ −50 days. A trimmed average or median is fine. However, your eyes are your best bet.

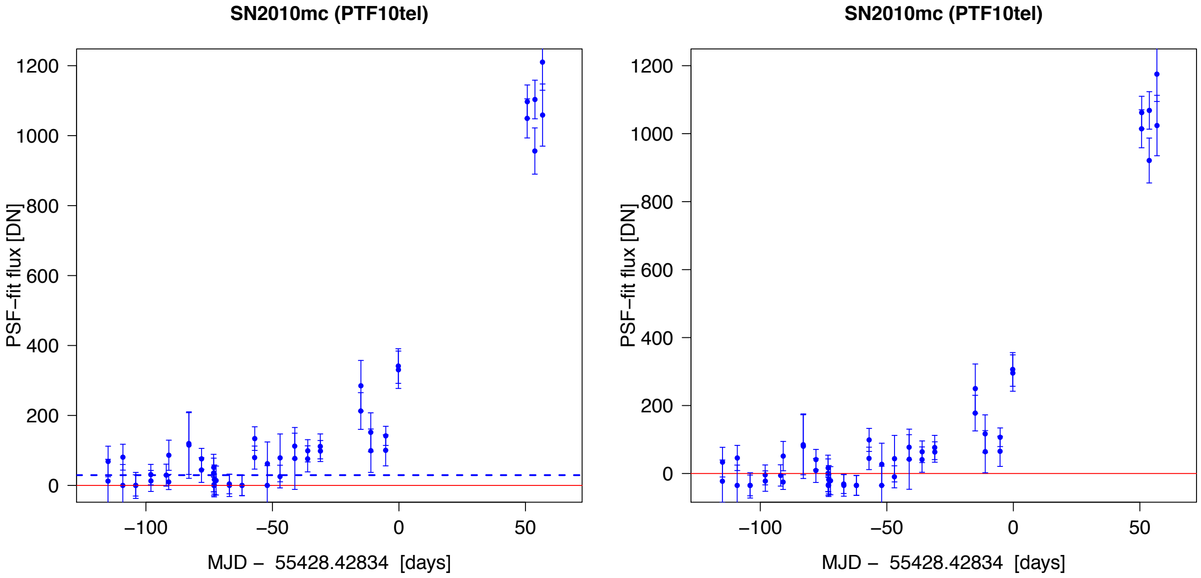

**Figure 1:** PSF-fit flux in DN ("flux" column) as a function of MJD. **LEFT:** uncorrected baseline where red line is the zero line and blue line is the baseline level estimated from a median of all fluxes at relative MJD < -50 days. **RIGHT:** corrected flux time-series with a constant baseline of 35 DN removed from all the measurements.

## 4 Uncertainty validation and possible rescaling

There is no guarantee the 1-σ uncertainties in the forced PSF-fit fluxes ("sigflux" values) are correct, or at least plausible with repeatability (*frequentist*) arguments. These are based on propagating semi-empirical models of the statistical (random) noise fluctuations expected in the detector-pixels through the difference- (and reference-) image pipelines. These models do not account for possible systematics, for example, incorrect PSF-template estimation, photometric zero-point calibrations (or refinements), astrometric calibrations (determining PSF-placement), and/or user-error in the supplied target positions.

There are three methods to validate these uncertainties, and correct them if needed. The last of these methods allows for a final coarse check, with possible tweaks. These methods are still crude, but they're simple and will get you closer to reality, if that matters. This check will also convey fidelity in your results. Personally, I would err on the conservative side and have my uncertainties slightly overestimated than underestimated. The latter leads to "overconfidence". You get the idea.

### Method I:

The first method entails examining the "chi" values in the forced photometry table. These represent the ratio of the RMS in the measured residuals in the PSF-fit to the average of all pixel uncertainties within the PSF-fitting area (default radius is 3 pixels). In effect, this metric is similar to the square-root of the classical reduced $\chi^2$ measure. Figure 2 shows an example of "chi" versus relative MJD for the same transient in Figure 1. The red line represents the desired value of one, indicating the uncertainties are plausible and consistent with fluctuations in the data, but not necessarily with all systematics. In this example, <chi> ~ 1.12, indicating the "sigflux" (DN) values will need to be inflated by a factor of 1.12. Therefore, as a rule of thumb:

sigflux(*corrected*) ≈ <chi> * sigflux(*raw*),

where <chi> represents either a trimmed average or median of the "chi" values within *some selected time range*. At the time of writing, we recommended that only historical measurements, i.e., prior to any explosive event be used to estimate the scaling factor <chi>. These can be the same measurements used to derive the baseline correction in Section 3. This scaling factor can then be applied to all the "sigflux(*raw*)" values in your table. It has not yet been verified, but the presence of a spatially non-uniform source-signal (e.g., when the transient is bright) could bias the chi estimates to high values, so don't be surprised to see flux-dependent chi-values. It is not yet clear if this is due to "bad" PSF-fitting or some deficiency underlying the chi computation. It is after all, an approximation to the true reduced $\chi^2$ metric.

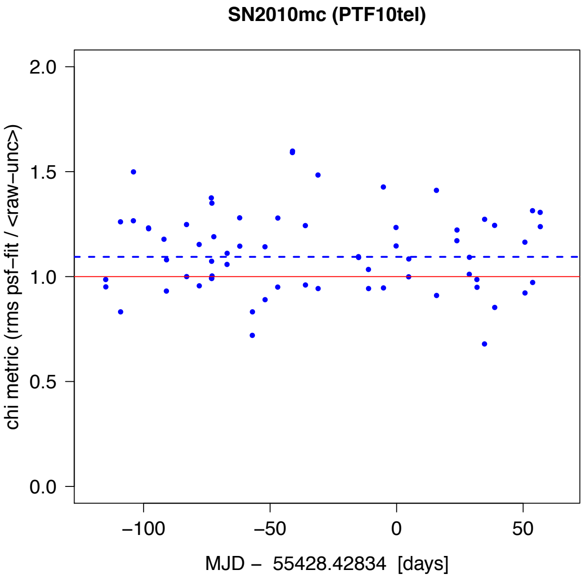

**Figure 2:** Distribution of "chi" values that can be used to validate and rescale the photometric uncertainties. Red line represents the desired value of one. Blue dashed line is the median value ~1.12, implying the measurement uncertainties ("sigflux" values) are underestimated by ~12%. See text for details.

### Method II:

The second method is probably more clear-cut, i.e., less subject to internal assumptions on how the "chi" metric is computed. This examines the variance in the historical flux (DN) measurements about a fixed (stationary) baseline. Again, these can be the same measurements used to derive the baseline correction in Section 3. One would expect the standard-deviation of these historical measurements to be consistent with the overall (e.g., mean or median) "sigflux" values. This assumes it's a stationary process, sampling presumably the same underlying "empty-sky" signal, so why not? Therefore, for transients at least, your global scaling factor would be the ratio:

$s = StdDev(\mathrm{flux}[t_{min} < t < t_{max}]) \,/\, < \mathrm{sigflux}[t_{min} < t < t_{max}] >$,

where < sigflux[…] > is ether a trimmed average or median of the sigflux values within the same time-range. You can then correct all the "sigflux(*raw*)" values in your table by multiplying by *s*. This method won't work for continuous variables (periodic or aperiodic) because of the non-stationarity in source flux. In this case, you may need to resort to light curve fitting (e.g., via Fourier series) and use the standard-deviation in the fit-residuals as a proxy for the photometric uncertainty, perhaps binned by flux (crudely speaking). Alternately, you could use the reduced $\sqrt{\chi^2}$ value from this fitting as the global scaling factor. This may sound like overkill. See the next paragraph for yet another ballpark method on how to quickly validate the "sigflux" values for the special case of flux-variables.

**Method III (final coarse check):**

For the astute reader, you may have noticed that methods I and II above (in the case of transients) use a scaling factor calibrated exclusively from the historical *zero-flux* "empty-sky" measurements. This will get you in the right ballpark for uncertainties on relatively faint fluxes, i.e., mostly at $R_{PTF} >\sim 17.5$ mag, but it says nothing about whether the photon-noise contribution for bright-sources was estimated correctly. I.e., whether the detector electronic-gain was properly propagated through all the processing and internal recalibration steps, including reference-image construction. Therefore, you may also want to check if the "sigflux" values in your table are more-or-less expected for the corresponding source-fluxes. This will be applicable to flux-measurements that are relatively bright in your table, e.g., for transients or flux-variables at/near their peaks. This check can be accomplished (again crudely) using a prior-calibrated photometric-repeatability (stack-RMS) versus magnitude plot. An example is shown in Figure 3. This plot was generated from PSF-fit photometry on difference images using 26,385 targets selected from the reference image, each consisting of >150 epochs. The reference image was made from 20 good-seeing frames, and is typical of most archived reference images. Figure 3 captures systematics from the image-differencing process, e.g., residuals from image misalignments, non-optimal photometric zero-point matching, PSF-matching, and/or systematics from the forced photometry step, e.g., epoch-based PSF-calibrations, PSF-centroiding, etc. When they occur, these systematics will have a greater impact for bright sources.

To use Figure 3, you will first need to convert your "flux" and "sigflux" (DN) values to magnitude equivalents (see Section 8). Of course this only makes sense for flux > 0 measurements, i.e., in the bright flux (high S/N) regime. Uncertainty validation in the low flux regime can be appropriately handled using methods I or II above. Figure 3 can be used as an independent check after any corrections from the above two methods were applied. For a given magnitude in hand (after converting from DN), you would read off the approximate "expected" $\sigma_{\mathrm{mag}}$ range from Figure 3 (i.e., spanned by the top and bottom green boundaries) and then check if your forced photometry $\sigma_{\mathrm{mag}}$ value ($\approx 1.0857\sigma_{\mathrm{DN}}/\mathrm{flux}_{\mathrm{DN}}$) falls within this range. If so, then all is consistent and you're good to go. If not, you could derive a scaling factor for the “sigflux” ($\sigma_{\mathrm{DN}}$) values in your table such that your $\sigma_{\mathrm{mag}}$ values (for the corresponding magnitudes) are consistent with the densest regions of the locus in Figure 3.

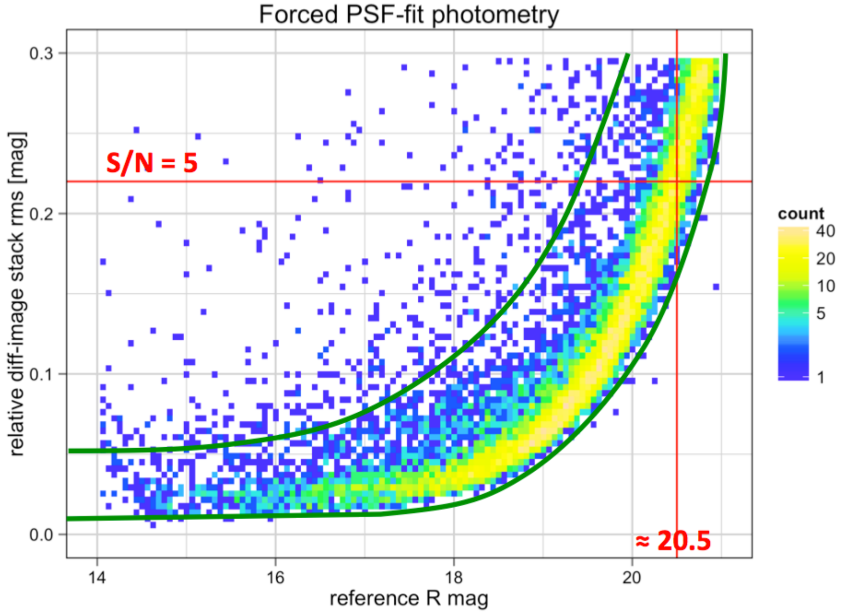

**Figure 3:** 1-sigma uncertainty in magnitude versus magnitude derived from the flux-RMS in forced PSF-fit photometry on >150 difference-image epochs at 26,385 target positions. The green lines define an approximate locus for the overall dependence. This locus can be used as a guide to check and possibly rescale the “sigflux” values in your table. See text for details.

## 5 Sanity checks using ancillary forced Aperture Photometry

So, when should you consider using the forced aperture-photometry measurements (the “fluxap” and “sigfluxap” columns)? One use case would be to replace any erroneous or missing PSF-fit photometry measurement at any epoch with the aperture measurement. Note that these aperture measurements are subject to the same caveats and potential corrections as in PSF-fit photometry (Sections 3 and 4). Furthermore, the aperture measurements have their own caveats: contamination by bad pixels (see the “nbadap” column), cosmic rays, saturation, source-confusion etc. PSF-fit photometry is relatively immune to these effects, including saturation if the saturated core of a source is smaller than the PSF-fitting area (default = 3 pixel radius).

Another more important issue with the aperture measurements is that these use a *fixed* aperture across all input epochs. The aperture is not adapted to the variable seeing and no curve-of-growth (or aperture) correction is applied. Currently, the default aperture radius is 6 pixels. This may not be enough, particularly when the seeing gets above ~ 4 pixels FWHM. Missed flux from extended PSF wings can lead to flux-deficits of ~ 1% or more. In future, we may make the aperture size user-specifiable. Be aware that larger apertures will lead to increased noise and contamination from bad pixels and/or source crowding. Therefore, proceed with caution when using the forced aperture-photometry measurements.

Aside from potentially using the aperture measurements as surrogates for bad/missing PSF-fit measurements (bearing in might the caveats and limitations from above; particularly missing flux when the seeing is "bad" or excess flux when source-confusion is high), you could also use them as an overall check of the PSF-fit photometry. For example, you could examine a plot of "1 – (flux / fluxap)" versus flux (DN) or mag to see if there is any global bias for flux estimates above zero. If so, it could mean bad PSF-placement, i.e., your supplied R.A., Dec. position could be in error with respect to that predicted by the difference-image astrometric solutions. The aperture measurements will be more immune to astrometric errors (i.e., aperture placement). It could also point to a global systematic in the calibrated PSF-templates for each exposure. Again, beware the caveats associated with the aperture measurements – these could be causing any bias you see. What shall you do if you see a significant difference between "flux" and "fluxap"? My suggestion is to correct (rescale) *all* the PSF-fit flux measurements for consistency. Again, proceed with caution and be aware of all the limitations associated with the forced-aperture measurements.

## 6 Other quality checks for PSF-fit photometry: the "sharp" diagnostic

As mentioned, the PSF-fit photometry is sensitive to the accuracy of the individual PSF-templates that are derived per exposure image. These are estimated using an automated process per exposure by selecting >~ 200 bright "isolated" stars in the PSF-matched (kernel-convolved) images prior to differencing. Functionality in the DAOPhot II package is used to perform this task. The PSF spatial variation is modeled using a linear function. Even though the process has proven to be robust, there's no guarantee the PSF estimates will always be perfect, particularly in regions with a high source density and/or complex background. A metric that could be used to assess whether the internally-derived PSF is a "good" match to the actual source being measured is the "sharp" diagnostic. This is a column with the same name in the forced-photometry table. In the end, all transients are expected to be unresolved and PSF-like, so deviations from this, for example from bad image-differencing, a bad PSF-matching kernel, or contamination by glitches and cosmic rays will render the PSF-fit bad and bias the flux estimate for the source in question. The "sharp" metric quantifies this deviation in a relative sense and hence can be used to assess PSF-fitted flux accuracy. Sharp is proportional to the difference: $\mathrm{FWMH}^2(\mathrm{obs}) - \mathrm{FWHM}^2(\mathrm{PSF\ template})$, where the first term is the square of the FWHM of the object inferred from the best fitting 2-D Gaussian, and the second term uses the FWHM of the PSF-template model interpolated at the object location. For *sources with a relatively high signal-to-noise ratio* (say >~ 7), abnormally low or high values of "sharp" can be used to flag or omit highly discrepant PSF-fit measurements. The various limits are defined as:

$$\mathrm{sharp}\begin{cases} << 0 \Rightarrow \text{cosmic ray, pixel spike, glitch} \\ \approx 0 \Rightarrow \text{source is "PSF-like", yea!} \\ >> 0 \Rightarrow \text{source has extended profile} \end{cases}$$

Again, the sharp metric only makes sense for sources with a relatively high signal-to-noise ratio. Figure 4 shows an example of "sharp" as a function of $R_{\mathrm{PTF}}$ magnitude for an iPTF *reference* image. This is expected to be applicable to difference-images as well. In this example, measurements with |sharp| > 0.4 should be scrutinized (visually on the images) and flagged or omitted from your light curve. Alternatively, you could keep their flux measurements but inflate their uncertainties ("sigflux" values) for consistency with other neighboring (presumably unbiased) light curve measurements. Easier said than done!

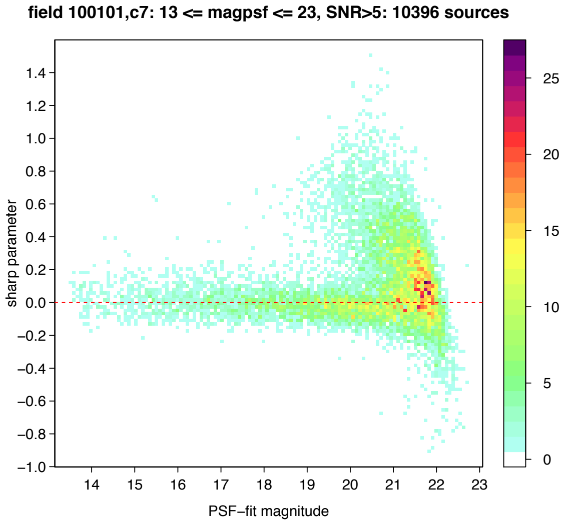

**Figure 4:** "sharp" metric as a function of R<sub>PTF</sub> PSF-fit magnitude for a *reference* image. This is expected to carry over to the noisier difference-images, where the 5-σ magnitude limit will more likely be R<sub>PTF</sub> ~ 20.5 mag. See text for details.

## 7 Computing flux Upper Limits for non-detections

It is assumed you have validated the photometric uncertainties ("sigflux" values) and corrected them if needed (Section 4). If for example your uncertainties are underestimated, your significance levels and derived upper limits will be wrong. "5-sigma" only has meaning if "sigma" is correct, i.e., if sigma is plausible given all noise sources that *could have* corrupted the "truth". As mentioned above, we suggest you err on the "more uncertain" (conservative) side when unsure if your uncertainties are correct. I.e., slight overestimation is safer for science in general, although one could argue there's no excuse for not getting them right in the first place! However, life isn't easy when hard-to-model systematics are at play.

Two subjective questions often asked are:

(i) what signal-to-noise threshold “SNT” should I assume for declaring a measurement a “non-detection” so it can be assigned an upper-limit?
and

(ii) what signal-to-noise value “SNU” should I use when computing a “SNU-sigma” flux upper-limit for a non-detection from its uncertainty (sigma) estimate?

Without going into details, here’s a choice: SNT = 3 and SNU = 5. If you want the details on why I picked these values for the simple case of Gaussian-distributed noise, see: *http://web.ipac.caltech.edu/staff/fmasci/home/mystats/UpperLimits_FM2011.pdf*
The choice is based on “detection probability”: a hypothetical source with flux equal to the upper-limit value (say 5-sigma) will have been detected above 3-sigma (and assigned a confidence interval based on sigma) with a probability of ~ 98%. That’s quite secure and shouldn’t cause any arguments, I hope.

## 8 Putting it all together: conversion to magnitudes

Now for the fun step – when you get to see the fruits of your labors. The pseudo-code conditional below makes use the following thresholds and parameters: “SNT” and “SNU” from Section 7; ZP (magnitude zero-point from the “zpmag” column in the forced-photometry table – see below); the possibly corrected PSF-fitted “flux” (Sections 3 & 5) and accompanying corrected “sigflux” (Section 4), both in DN; and from these, the SNR (≡ flux/sigflux).

```
if( SNR > SNT )
  # we have a “confident” detection, compute and plot mag with error bar:
  mag = ZP – 2.5*$\log_{10}$[flux]
  $\sigma_{mag}$ = 1.0857*sigflux / flux
else
  # compute flux upper limit (mag lower limit) and plot as arrow or triangle:
  mag = ZP – 2.5*$\log_{10}$[SNU*sigflux]
```

At the time of writing, the “zpmag” (ZP) values provided in the table are computed using the reference-image SExtractor catalog MAG_AUTO measurements using as many “clean” stars as possible within 14.5 ≤ *R* ≤ 19.0. These are matched and differenced with fixed “big-aperture” instrumental magnitudes on a per-frame (epoch) basis. The ZP is then computed as a median of these differences. The MAG_AUTO magnitudes are the only measurements that are (indirectly) tied to SDSS absolute photometry in our system, i.e., as inherited from the IPAC frame-processing pipeline during reference-image construction.

The zpmag values enable one to convert the instrumental fluxes from either PSF-fitting or “big-aperture” photometry to calibrated magnitudes (assuming both catch the *same total* instrumental flux). By “big-aperture”, we mean apertures with radii >~ 6 arcsec, where 6 arcsec is the default fixed aperture radius used to compute the “fluxap” values in the table. One caveat is that if the seeing (“FWHMSEX” column) is bad, e.g., >~ 4.5 arcsec, a 6 arcsec aperture will not catch the total instrumental flux and hence will not be consistent with the flux from PSF-fitting. In this case, the "zpmag" value will not be applicable to the "fluxap" value.

The "zprms" values in the forced-photometry table provide an estimate of the systematic uncertainty in the absolute photometric calibration. These are computed using a robust RMS of the absolute-to-instrumental magnitude differences. If any "zpmag, zprms" values happen to equal "27, 0" for any epoch, it means a calibration could not be performed (e.g., not enough good stars). In this case, we advise using an average or median of the zpmag values (that are ≠ 27) from other epochs in the light curve.

Figure 5 shows an example after applying the above pseudo-code to the corrected flux, sigflux [DN] values from Figure 1. Here, we assumed SNT = 2 (not 3) to get more points with error bars to show up, and SNU = 5. Note that the magnitude upper limits are brighter (more conservative) by design.

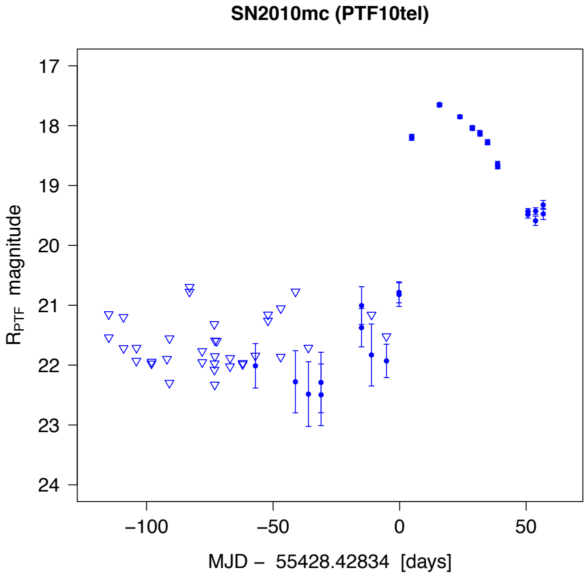

**Figure 5:** example light curve after converting the flux [DN] measurements from Figure 1 (right panel) to upper-limits (triangles) and detections (circles with error bars) to magnitudes.

## 9 AC–to–DC Photometry for variables (adding a non-zero flux level)

Since transients are commonly associated with sudden events where a signal suddenly appears out of the noise then fades, i.e., emerging from locations where there is (usually) no detectable prior signal, one need not worry about adding a DC offset. The goal (usually) is to analyze the energy released during the explosive phase (in excess of any pre-existing signal, if any). For continuous (a)periodic variables, or transients associated with pre-quiescent sources with detectable signal, or sources that could suddenly disappear, one may want to characterize the total (DC) signal versus time. Here, you would need to obtain an estimate of the flux of your source from the *same* reference-image used to generate the difference images in your forced photometry request. Note that no source may be visible on the reference image. Therefore, for your convenience, we also compute forced PSF-fit photometry at the same R.A., Dec. position on the reference-image and write this to the output light curve table (Appendix I). This information is given immediately after the epochal light curve measurements for the specific source ID, for example:

```
\ Reference-image PSF-fit metrics for id 1: flux_ref, sigflux_ref, chi_ref, sharp_ref = 1677.8 DN,
49.7257 DN, 1.623, -0.076
```

If forced-photometry on the reference image could not be performed (e.g., due to a missing PSF in the archive or some other processing glitch), the values above will be –99.999.

If interested in the DC light curve (modulated by AC variations inferred from image-differencing), below is the new pseudo-code. Inputs related to the AC component are “flux” and “sigflux” (DN) and ZP (“zpmag”) from the forced-photometry table, and “flux_ref” and its uncertainty “sigflux_ref” (both in DN) at the end of each light curve. First, the signal-to-noise ratio, SNR, is redefined as:

$\mathrm{SNR} = (\mathrm{flux} + \mathrm{flux\_ref}) / \mathrm{sigflux\_DC}$,

where
$\mathrm{sigflux\_DC} = \sqrt{(\mathrm{sigflux}^2 - \mathrm{sigflux\_ref}^2)}$ if $\mathrm{sigflux} > \mathrm{sigflux\_ref}$,

otherwise,
$\mathrm{sigflux\_DC} = \sqrt{(\mathrm{sigflux}^2 + \mathrm{sigflux\_ref}^2)}$.

You may be asking: why are the variances subtracted from each other and not added under the first square-root? This follows from the assumption that “flux_ref” and “sigflux_ref” are measured on the *same* reference image used to generate the “single-exposure – reference” differences: “flux” and “sigflux”. I.e., the noise is expected to be anti-correlated. This also assumes the reference and single-exposure images are independent, i.e., a sufficient number of exposures were used to create the reference image that correlations with any of the individual input exposures are minimal.

If for whatever reason *any* of the following are *true*:
(i) sigflux ≤ sigflux_ref, e.g., due to some inadvertent detail when propagating the image-noise,
(ii) an entirely different reference-image were used to infer “flux_ref” and “sigflux_ref”,
(iii) the sigflux_DC estimates appear unreasonably low compared to the ballpark expectations of Figure 3,
we suggest adding the noise contributions in quadrature (second square-root expression above). This is a conservative estimate.

With SNR defined above and SNT, SNU defined in Section 7, here's the DC light curve generation logic:

```
if( SNR > SNT )
  # we have a "confident" detection, compute and plot mag with error bar:
  mag = ZP – 2.5*log$_{10}$[flux + flux_ref]
  $\sigma_{mag}$ = 1.0857*sigflux_DC / (flux + flux_ref)
else
  # compute flux upper limit and plot as arrow or triangle:
  mag = ZP – 2.5*log$_{10}$[SNU*sigflux_DC]
```

You may notice that this logic reduces to the pure AC light curve generation logic in Section 8 when flux_ref = 0 and sigflux_ref = 0 (or equivalently, sigflux_DC = sigflux).

## 10 Going deeper: combining single-epoch measurements

If you're feeling ambitious and want to make your light curve measurements more statistically significant and/or want tighter constrains on flux upper-limits, you can attempt to combine your flux (DN) measurements within carefully selected time-windows using some optimal method. This procedure assumes of course your single-epoch uncertainties have been validated and corrected if necessary (Section 4). One method is to *assume* the underlying source signal *is stationary (flat)* within a time-window and collapse the single-epoch ($flux_i$) measurements therein using an inverse-variance weighted average:

$$flux_{comb} = \frac{\sum_i w_i flux_i}{\sum_i w_i} \quad \text{where} \quad w_i = \frac{1}{sigflux_i^2},$$

with uncertainty:

$$sigflux_{comb} = \left[ \sum_i w_i \right]^{-1/2}.$$

The last line reduces to $sigflux_{comb} \approx sigflux/\sqrt{n}$ if one assumes the $n$ single-epoch measurements in a window have approximately the same uncertainty equal to some constant *sigflux*. Therefore, the improvement in signal-to-noise assuming uncorrelated inputs cannot be more than a factor of $\approx \sqrt{n}$. The important assumption here is that the underlying source emission is constant. It may appear constant within measurement error, but when many measurements are available, you can attempt to bin them in different ways to attempt to tease out possible hidden trends in the signal. A moving (weighted) average might also work, but be careful not to smooth-out real variations. There's also a huge collection of methods on local-polynomial regression fitting. Feel free to experiment.

You can also attempt to collapse the measurements within windows by fitting a prior model of flux versus time, i.e., if you have prior (or contextual) knowledge that previous deeper observations revealed that your source exhibited a linear or non-linear trend. More information the better! The advantages of combining measurements in source-space across epochs as opposed to co-adding entire images in time-ordered slices are: (i) speed and (ii) greater flexibility in the combination method.

After having window-combined some of the flux (DN) measurements, you can convert to magnitudes and assign upper-limits using exactly the same light curve generation logic presented in Section 8 (for transients above a zero-flux level) or Section 9 (for variables). The important thing here is that the AC single-epoch "flux, sigflux" measurements for transients, or the "flux+flux_ref, sigflux_DC" measurements for variables are replaced with $flux_{comb}$ , $sigflux_{comb}$ respectively. Figure 6 shows an example of collapsing the flux (DN) measurements within specially selected time-windows using weighted averaging, then converting to magnitudes.

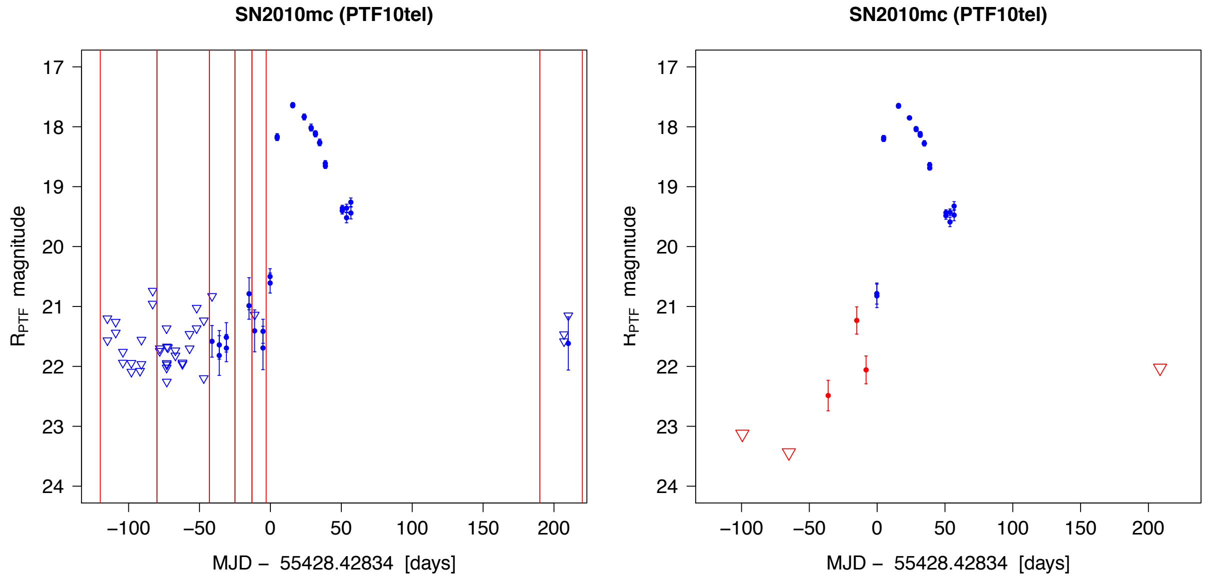
SN2010mc (PTF10tel)

SN2010mc (PTF10tel)

**Figure 6: LEFT:** same example light curve as above (but over a longer time span) showing the time-window boundaries within which measurements are combined. Blue triangles are single-epoch upper limits and blue circles with error bars are detections. **RIGHT:** result of combining measurements using inverse-variance weighted averaging of the *fluxes* (in DN) within each window on the left. Red triangles are combined-flux upper limits; red circles with error bars are combined-flux detections; blue points are the original single-exposure measurements.

# 11 Appendices

## I. Example of a forced photometry table

**Note:** the data rows in the example table below are wrapped and truncated. This table is in the traditional IPAC-style format, consisting of a header, column descriptors, data types, and units for each column. This table may contain multiple light curves corresponding to different *sourceid* entries, for example, if a list of targets for the same field, chip, and filter were submitted.

**Generic table filename:** forcepsffitdiff_d<fieldID>_f<filterID>_c<chipID>.out

```
\ Forced PSF-fit and aperture photometry results
\ Generated by forcepsffitdiff.pl v3.0, 2015-08-10 at 10:55:58
\ Input PSF-fitting radius = 3 pixels
\ Magnitude Zero-Point & RMS: see epoch-dependent zpmag,zprms columns
\ Input (fixed) aperture radius = 6 pixels
\ Input (fixed) inner radius of sky annulus = 12 pixels
\ Input (fixed) outer radius of sky annulus = 16 pixels
\ Input (fixed) aperture correction = 0 magnitudes
\ Input correlated-noise correction factor for aperture photometry = 1
\ Number of target positions = 1
\ Column definitions:
\ sourceid = source ID based on internal counter
\ xpos, ypos = source centroid position in x,y system of difference image
\ ra, dec = corresponding J2000 equatorial coordinates
\ flux, sigflux = PSF-fit flux and 1-sigma uncertainty; use these estimates
\        when generating light curves; values of 99999999,-99.999 indicate
\        no measurement was possible.
\ fluxap, sigfluxap = aperture flux and 1-sigma uncertainty using a _fixed_
\        aperture. Provides a crude sanity check of PSF photometry. Will
\        be affected by bad pixels (see nbadap column). Note: aperture
\        correction (if applied) may not match variable seeing. You may
\        want to increase aperture radius to make this more immune to
\        variable seeing. This will be at the expense of an increase in the
\        measurement noise.
\ nbadap = number of bad and saturated pixels in measurement aperture;
\        99999999 => aperture measurement not possible (e.g., too close to edge)
\ snr = signal-to-noise ratio in PSF-fit flux estimate
\ chi = robust estimate of ratio: RMS in PSF-fit residuals /
\                                 expected RMS using uncertainties
\ sharp = fwhm_obs^2 - fwhm_PSF^2;
\       ~ 0  => perfect (source is PSF-like);
\       >> 0 => extended source or bad kernel solution;
\       << 0 => cosmic ray, glitch, or bad kernel solution
\ MJD, HJD = Modified and Heliocentric Julian Dates at start of exposure
\ FWHMSEX = Effective FWHM (seeing) of original science frame
\
\ Reference-image rfId, refimage = 172008,
/ptf/pos/sbx1/refims/d22100/f2/c4/p12/v1/PTF_d022100_f02_c04_u000172008_p12_refimg.fits
\ Reference-image PSF-file rfaId, rawpsf = 631890,
/ptf/pos/sbx1/refims/d22100/f2/c4/p12/v1/PTF_d022100_f02_c04_u000172008_p12_daopsf.rpsf
\
|sourceid    |xpos        |ypos        |ra          |dec         |zpmag       |zprms       |flux
|sigflux           |snr         |chi         |sharp       |fluxap              |sigfluxap         |nbadap
|MJD                 |HJD                   |FWHMSEX                   |
|i           |r           |r           |r           |r           |r           |r           |r
|r                 |r           |r           |r           |r                   |r                 |i
|r                   |r                     |r                         |
|            |pix         |pix         |deg         |deg         |mag         |mag         |DN
|DN                |            |            |            |DN                  |DN                |
|days                |days                  |arcsec                    |
 1             387.38      2139.94      200.4898400   11.7357530  27.061        0.039               -34.1965000
43.2692000    -0.79        1.014       -2.982               85.4956512        142.7702269  0
56442.20184           2456442.70502                 2.35000
 1             405.07      2114.14      200.4898400   11.7357530  27.062        0.038               -69.0671000
33.1654000    -2.08        1.271        0.263             -117.6785583        128.1161394  0
56442.23682           2456442.73999                 1.75000
etc ...
\ Reference-image PSF-fit metrics for id 1: flux_ref, sigflux_ref, chi_ref, sharp_ref = 1677.8 DN, 49.7257 DN,
1.623, -0.076
\ -----------------------------------------------------------------------------------------------------------
```

## II. Further reading and references

This list represents a relatively unbiased compendium of tidbits, tricks, and other information you may find useful.

- Computing flux upper-limits for non-detections:
  *http://web.ipac.caltech.edu/staff/fmasci/home/mystats/UpperLimits_FM2011.pdf*
- The danger of statistical inference in magnitude space:
  *http://web.ipac.caltech.edu/staff/fmasci/home/mystats/logfluxbias.html*
- Noise-variance (and sigma) in magnitude space:
  *http://web.ipac.caltech.edu/staff/fmasci/home/mystats/NoiseVarMagSpace.pdf*
- A Study of the bias from inverse Poisson-variance weighting:
  *http://web.ipac.caltech.edu/staff/fmasci/home/mystats/poisson.html*
- Optimal Image Combination in the presence of variable seeing:
  *http://web.ipac.caltech.edu/staff/fmasci/home/mystats/ImCombineWseeing.pdf*
- A quick 'n dirty way to estimate point-source sensitivity from image data:
  *http://web.ipac.caltech.edu/staff/fmasci/home/mystats/QuickMagLimit.txt*
- Optimum aperture size for a Gaussian light-profile:
  *http://web.ipac.caltech.edu/staff/fmasci/home/mystats/GaussApRadius.pdf*
- Photometric uncertainty estimation assuming priors and correlated image-noise:
  *http://web.ipac.caltech.edu/staff/fmasci/home/mystats/ApPhotUncert_corr.pdf*
- Simple photometric uncertainty estimation without priors:
  *http://web.ipac.caltech.edu/staff/fmasci/home/mystats/ApPhotUncert.pdf*
- iPTF (and ZTF) Image Differencing & Extraction (PTFIDE) presentation – includes the latest improvements to PTFIDE:
  *http://web.ipac.caltech.edu/staff/fmasci/home/miscscience/masci_lsst_ztf_Nov2014.pdf*
- The original PTFIDE document – ancient but some sections still relevant:
  *http://web.ipac.caltech.edu/staff/fmasci/home/miscscience/ptfide-v4.0.pdf*